

CLAIMS

What is claimed is:

- 1 1. A method for replicating data from a master server to a slave server over
2 a network, the method comprising the steps of:
3 sending a packet of information from the master server to the
4 slave server, the information relating to a change in
5 the data stored on the master server and containing a
6 version number for the present state of the data;
7 allowing the slave server to determine whether the data on
8 the slave server has been updated to correspond to
9 the version number contained in the packet; and
10 requesting a delta be sent from the master server to the
11 slave server if the data on the slave server does not
12 correspond to the version number contained in the
13 packet, the delta containing information needed to
14 update the slave server.
- 1 2. A method according to claim 1, further comprising:
2 storing an original copy of the data on the master server.
- 1 3. A method according to claim 1, further comprising:
2 persistently caching the data on a local disk for each slave server.

1 4. A method according to claim 1, further comprising:
2 determining a unique version number for the current state of the
3 data on the master server if the data has changed.

1 5. A method for replicating data from a master server to a slave server over
2 a network, the method comprising the steps of:

3 sending a version number from the master server to the
4 slave server, the version number relating to the
5 present state of the data stored on the master server;
6 allowing the slave server to determine whether the slave
7 server has been updated to reflect the present state of
8 the data corresponding to the version number sent
9 from the master server; and
10 requesting a delta be sent from the master server to the
11 slave server if the slave server does not correspond to
12 the version number sent by the master, the delta
13 containing information needed to update the slave
14 server.

1 6. A method according to claim 5, further comprising:
2 sending the delta from the master server to the slave server.

1 7. A method according to claim 5, further comprising:

2 committing the delta to the slave server.

1 8. A method according to claim 5, further comprising:

2 updating the version number of the slave server after committing the
3 delta.

1 9. A method according to claim 5, further comprising:

2 periodically sending the version number from the master server to
3 a slave server.

1 10. A method according to claim 5, further comprising:

2 sending the version number to a slave server until the slave server
3 acknowledges receipt of the version number.

1 11. A method according to claim 5, further comprising:

2 including data with the version number that is necessary to update
3 a slave server.

1 12. A method according to claim 11, further comprising:

2 committing the data necessary to update the slave server as soon
3 as it is received.

1 13. A method according to claim 5, further comprising:

2 determining the scope of the delta before sending it from the master
3 server.

1 14. A method for replicating data over a network including a master
2 server and at least one slave server, the method comprising the
3 steps of:

4 sending a packet of information from a master server to each
5 slave server on the network, the information relating to
6 a change in the data stored on the master server and
7 containing a current version number for the present
8 state of the data, the information further relating to
9 previous changes in the data and a version number
10 for each previous change;

11 allowing each slave server to determine whether the slave
12 server has been updated to correspond to the current
13 version number;

14 allowing each slave server to commit the information if the
15 slave server has not missed a previous change; and

16 allowing each slave server having missed a previous change
17 to request that previous change be sent from the
18 master server to the slave server before the slave
19 server commits the packet of information.

1 15. A method according to claim 14, further comprising:

2 committing the packet of information to a slave server.

1 16. A method according to claim 14, further comprising:

2 aborting the commit of the packet of information if a slave server
3 cannot commit the update.

1 17. A method according to claim 14, further comprising:

2 determining the scope of the delta before sending it from the master
3 server.

1 18. A method according to claim 14, further comprising:

2 including the scope of each the previous changes in the delta.

1 19. A method for replicating data over a network including a master server
2 and at least one slave server, the method comprising the steps of:

3 sending a packet of information from a master server to each
4 slave server on the network, the information relating to
5 a change in the data stored on the master server and
6 containing a prior version number for the prior state
7 and a new version number for the new state of the
8 data, the information further relating to previous
9 changes in the data and a previous version number

10 for each previous change;
11 allowing each slave server to determine whether the data on
12 the slave server corresponds to the prior version
13 number contained in the packet;
14 allowing each slave server to commit the packet of
15 information if the data on the slave server corresponds
16 to the prior version number contained in the packet,
17 the commit also updating the version of the slave
18 server to the new version number; and
19 allowing each slave server not corresponding to the prior
20 version number to request that a delta be sent from
21 the master server containing the information
22 necessary to update the slave to the prior version
23 number before the slave server commits the packet of
24 information.

1 20. A method for replicating data over a network including a master server
2 and at least one slave server, the method comprising the steps of:
3 sending a packet of information from a master server to each
4 slave server on the network, the information relating to
5 a change in the data stored on the master server and
6 containing a version number for the prior state and a
7 version number for the new state of the data, the

8 information further relating to previous changes in the
9 data and a version number for each previous change;
10 allowing each slave server to determine whether the data on
11 the slave server corresponds to the prior version
12 number contained in the packet;
13 allowing each slave server to commit the packet of
14 information if the data on the slave server corresponds
15 to the prior version number contained in the packet,
16 the commit also updating the version of the slave
17 server to the new version number; and
18 allowing each slave server not corresponding to the prior
19 version number to request that a delta be sent from
20 the master server containing the information
21 necessary to update the slave to the new version
22 number.

1 21. A method for replicating data from a master server to at least one slave
2 server over a network, the method comprising the steps of:
3 sending a packet of information from the master server to a
4 slave server, the information relating to a change in
5 the data stored on the master server and containing a
6 version number for the present state of the data;
7 receiving the packet of information to a slave server;

8 allowing the slave server to determine whether the slave
9 server has been updated to correspond to the version
10 number contained in the packet, and to further
11 determine whether the slave server can process the
12 packet of information if needed to update to
13 correspond to the version number contained in the
14 packet;
15 sending a signal from the slave server to the master server,
16 the signal indicating whether the slave server needs to
17 be updated and whether the slave server can process
18 the update; and
19 sending a response signal from the master server to the
20 slave server indicating whether the slave server
21 should commit to the information contained in the
22 packet; and
23 committing the packet of information to the slave server if so
24 indicated by the response signal.

1 22. A method according to claim 21, further comprising:
2 determining whether each of the at least one slave server can
3 commit the data.

1 23. A method according to claim 21, further comprising:

2 determining whether each of the at least one slave server has sent
3 a response back to the master server.

1 24. A method according to claim 21, further comprising:
2 determining whether any of the at least one slave server can commit
3 the data.

1 25. A method according to claim 21, further comprising:
2 committing the data only if each of the at least one slave server can
3 process the commit.

1 26. A method according to claim 21, further comprising:
2 aborting the data only if any of the at least one slave server cannot
3 process the commit.

1 27. A method according to claim 21, further comprising:
2 committing the data to those slaves that are able to process the
3 commit.

1 28. A method according to claim 21, further comprising:
2 multicasting the update to any of the at least one slave server that
3 were not able to process the commit.

1 29. A method according to claim 21, further comprising:
2 heartbeating the new version number to any of the at least one
3 slave server that were not able to process the commit.

1 30. A method according to claim 21, further comprising:
2 requesting a delta be sent to a slave server that was not able to
3 process the commit.

1 31. A method for replicating data over a network, the method comprising
2 the steps of:

3 (a) determining whether the replication should be accomplished in
4 a one or two phase method;

5 (b) sending replication information determined to be accomplished
6 in a one phase method by:

7 sending a packet of information from the master server to the
8 slave server, the information relating to a change in
9 the data stored on the master server and containing a
10 version number for the present state of the data;

11 receiving the packet of information to a slave server;

12 allowing the slave server to determine whether the data on

13 the slave server has been updated to correspond to

14 the version number; and

15 requesting a delta be sent from the master server to the

16 slave server if the slave server does not correspond to
17 the version number, the delta containing information
18 needed to update the slave server;
19 (c) sending replication information determined to be accomplished
20 in a two phase method by:
21 sending a packet of information from the master server to the
22 slave server, the information relating to a change in
23 the data stored on the master server and containing a
24 version number for the present state of the data;
25 allowing the slave server to determine whether the slave
26 server has been updated to correspond to the version
27 number, and to further determine whether the slave
28 server can process the packet of information;
29 sending a signal from the slave server to the master server
30 indicating whether the slave server needs to be
31 updated and whether the slave server can process the
32 packet of information;
33 sending a response signal from the master server to the
34 slave server indicating whether the slave server
35 should commit to the packet of information; and
36 committing the packet of information to the slave server if so
37 indicated by the response signal.

32. A method for replicating data over a network, the method comprising the steps of:

(a) determining whether replication should be accomplished in a one or two phase method;

(b) sending data to be replicated in a one phase method by:

sending a version number for the current state of the data from a master server to a slave server;

requesting a delta be sent from the master server to the slave server if the data on the slave server does not correspond to the version number; and

(c) sending data to be replicated in a two phase method by:

sending a packet of information from the master server to a slave server;

determining whether the slave server can process the packet of information; and

committing the packet of information to the slave server if the slave server can process the packet of information.

33. A method for replicating data from a master to a plurality of slaves on a network, the method comprising the steps of:

(a) determining whether replication should be accomplished in a one or two phase method;

(b) sending data to be replicated in a one phase method by:

6 sending a version number for the current state of the data
7 from the master to each slave; and
8 requesting a delta be sent from the master to each slave
9 containing data that does not correspond to the
10 version number;

11 (c) sending data to be replicated in a two phase method by:
12 sending a packet of information from the master to each
13 slave; and
14 committing the packet of information to the slaves if each of
15 the plurality of slaves can process the packet of
16 information.

1 34. A method for replicating data from a master to a plurality of slaves on
2 a network using one and two phase methods, the method comprising the
3 steps of:

4 (a) sending data to be replicated in a one phase method by sending
5 a version number for the current state of the data from the
6 master to each slave so that each slave may request a delta
7 to be sent from the master to the slave to update the data on
8 the slave; and

9 (b) sending data to be replicated in a two phase method by sending
10 a packet of information from the master to each slave, the
11 packet of information to be committed by each slave if every

12 slave is able to commit the packet of information.

1 35. A method for replicating data on a clustered network using one and two
2 phase methods, each network cluster containing a cluster master and at
3 least one cluster slave, the method comprising the steps of:

4 (a) sending data to be replicated in a one phase method by sending
5 a version number for the current state of the data from a first
6 cluster master to all other cluster masters so the other cluster
7 masters may each request a delta; and

8 (b) sending data to be replicated in a two phase method by sending
9 a packet of information from the first cluster master to each
10 other cluster master, the packet of information to be
11 committed by the other cluster masters if the other cluster
12 masters are able to commit the packet of information.

1 36. A method according to claim 35, further comprising:
2 sending the data from each cluster master to each cluster slave in
3 the cluster with that cluster master by a one-phase method.

1 37. A method according to claim 10, further comprising:
2 sending the data from each cluster master to each cluster slave in
3 the cluster with that cluster master by a two-phase method.

1 38. A computer-readable medium, comprising:

2 (a) means for sending a packet of information from a master server
3 to each slave server on the network, the information relating
4 to a change in the data stored on the master server and
5 containing a current version number for the present state of
6 the data, the information further relating to previous changes
7 in the data and a version number for each previous change;

8 (b) means for allowing each slave server to determine whether the
9 slave server has been updated to correspond to the current
10 version number;

11 (c) means for allowing each slave server to commit the information
12 if the slave server has not missed a previous change; and

13 (d) means for allowing each slave server having missed a previous
14 change to request that previous change be sent from the
15 master server to the slave server before the slave server
16 commits the packet of information.

1 39. A computer program product for execution by a server computer for
2 replicating data over a network, comprising:

3 (a) computer code for sending a packet of information from a master
4 server to each slave server on the network, the information
5 relating to a change in the data stored on the master server
6 and containing a current version number for the present state

7 of the data, the information further relating to previous
8 changes in the data and a version number for each previous
9 change;

10 (b) computer code for allowing each slave server to determine
11 whether the slave server has been updated to correspond to
12 the current version number;

13 (c) computer code for allowing each slave server to commit the
14 information if the slave server has not missed a previous
15 change; and

16 (d) computer code for allowing each slave server having missed a
17 previous change to request that previous change be sent
18 from the master server to the slave server before the slave
19 server commits the packet of information.

1 40. A system for replicating data over a network, comprising:

2 (a) means for sending a packet of information from a master
3 server to each slave server on the network, the
4 information relating to a change in the data stored on
5 the master server and containing a current version
6 number for the present state of the data, the
7 information further relating to previous changes in the
8 data and a version number for each previous change;

9 (b) means for allowing each slave server to determine whether the

10 slave server has been updated to correspond to the current
11 version number;
12 (c) means for allowing each slave server to commit the information
13 if the slave server has not missed a previous change; and
14 (d) means for allowing each slave server having missed a previous
15 change to request that previous change be sent from the
16 master server to the slave server before the slave server
17 commits the packet of information.

1 41. A computer system comprising:
2 a processor;
3 object code executed by said processor, said object code configured
4 to:
5 (a) send a packet of information from a master server to
6 each slave server on the network, the information
7 relating to a change in the data stored on the master
8 server and containing a current version number for the
9 present state of the data, the information further
10 relating to previous changes in the data and a version
11 number for each previous change;
12 (b) allow each slave server to determine whether the slave
13 server has been updated to correspond to the current
14 version number;

- 15 (c) allow each slave server to commit the information if the
16 slave server has not missed a previous change; and
17 (d) allow each slave server having missed a previous change
18 to request that previous change be sent from the
19 master server to the slave server before the slave
20 server commits the packet of information.